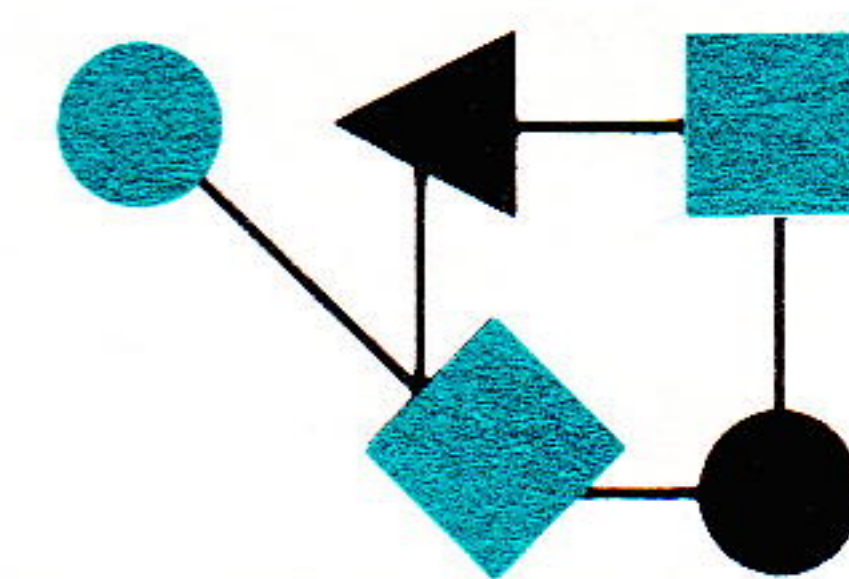


CONNECTIONS



The Interoperability Report

June 1987

Volume 1, No. 2

*ConneXions -
The Interoperability Report
tracks current and emerging
standards and technologies
within the computer and
communications industry.*

In this issue:

Network Management	2
NetBIOS update.....	2
Revised GOSIP.....	3
IETF reorganized.....	3
High-Level Management....	4
ISO OSI at Wisconsin.....	7
Profile: NSFNET.....	8
Padlipsky on Whitecaps.....	12

ConneXions is published by Advanced Computing Environments, 480 San Antonio Road, Suite 100, Mountain View, CA 94040, USA. Phone: 415-941-3399.

© 1987
Advanced Computing Environments.
Quotation with attribution encouraged.

ConneXions-The Interoperability Report
and the *ConneXions* masthead are
trademarks of Advanced Computing
Environments.

ISSN 0894-5926

From the Editor

One of the primary roles of *ConneXions* is to report events which affect the interoperability industry. Perhaps the most significant such event in recent weeks was the formation of a Network Management Working Group (see page 2). Several systems for dealing with network management are being developed. This issue presents an article by the authors of one such system, the High-Level Entity Management System which is being developed on the DARPA/NSF Internet.

We also feel that it is important to give you an overview of existing networks which are based on TCP/IP. This month we look at NSFNET and its component networks.

Much effort is currently going into implementation of OSI protocols, and in this issue we have an overview of the work being done at the University of Wisconsin. For those interested in learning more about OSI protocols and their implementation, Advanced Computing Environments is sponsoring an "ISO Development Seminar" in September 1987 in Monterey, California. The seminar will survey OSI standards with critical comments and implementation advice, together with reports on practical implementation experiences.

Speakers at the seminar will include Rob Hagens, Nancy Hall, Sue Lebeck, and Professors Lawrence Landweber and Marvin Solomon, all of the University of Wisconsin, as well as Dr. Marshall Rose of Northrop Research and Development Center. Dr. Rose has implemented several of the OSI upper layers and is author of the ISO Development Environment (ISODE). (See also "ISODE: Horizontal Integration in Networking," *ConneXions*, May 1987.)

Our guest editorial is by Michael A. Padlipsky. The name needs little introduction to the Internet community. New readers may think his style is a little unusual, but we hope you find him interesting reading nonetheless.

More vendors have joined the already large group who support TCP/IP and the underlying network protocols. **Adax Inc.**, **Mitek**, **Eastman Communications**, and **DevelCon** should be added to the list we printed in our Premiere Issue. We also should note that **Bellcore** is *not* a vendor of TCP/IP products at this time. Apologies for the error.

Network Management Working Group meeting held

On May 5, 1987, some 90 vendors, users and researchers met at Techmart in Santa Clara, California, to discuss the formation of a Network Management Working Group. There is currently a large interest in this arena, and several groups are working independently on management architectures. This meeting should be seen as an attempt to combine and coordinate this work, and achieve results in the short term which will benefit users.

Needed soon

Stan Ames and Lee LaBarre from the Mitre Corporation have stepped forward as initiators in trying to accomplish a multi-vendor architecture which can be deployed in products within a year. Stan first gave an overview of their objectives. In his opinion the group should:

- Write RFCs which define the architecture and the kinds of parameters which need to be standardized.
- Limit the scope; network management could easily turn into a massive effort. We need something in 6-12 months.
- Make sure we don't re-invent the wheel, by closely following the efforts of other groups (ANSI, IEEE, NBS, MAP/TOP, etc.) and working under the supervision of the Internet Engineering Task Force.

Groups can help each other

There were several presentations by various groups working in the area of network management. Although each of these groups has very specific environments in mind, there is a lot of common interest and it is hoped that they can coordinate their efforts.

On May 14th a "core group" of about 15 people met at Stanford to begin defining the framework in which the network management group will be working. The approach taken was the formation of several small working groups who will report back to the larger group from time to time. The next full meeting is currently targeted for the end of July in Washington, DC. We will keep you informed as this work progresses.

NetBIOS update

More feedback wanted

As we reported in our May issue, RFC 1001 and 1002 describing the NetBIOS service on top of TCP/IP have been issued. The authors have received a slow stream of feedback from the community. Specifically, there has been some discussion about how the NetBIOS work could be ported to the ISO environment. In addition, several companies are working on NetBIOS implementations. It is expected that the experience gained from these implementations coupled with the comments received will result in an updated set of RFCs by the end of this year. The authors are still actively soliciting input. If you have comments or questions, please send a message to Avnish Aggarwal: "mtxinu!excelan!avnish@ucbvax.berkeley.edu" or write:

Karl Auerbach
Epilogue Technology Corporation
P.O. Box 5432
Redwood City, CA 94063

New version of GOSIP published

Profile instead of Procurement Spec

On April 22, 1987, NBS published a revised version of GOSIP. This document contains a number of technical changes as compared with the December version, but the major change was in the scope and applicability section. GOSIP now stands for Government Open System Interconnection *Profile*, rather than Procurement Specification, and is accompanied by a Draft Federal Information Processing Standards Publication (FIPS) which (when approved) adopts GOSIP as a Federal Standard.

Does not exclude other technologies

Readers of the earlier document raised concern that government agencies would have to start procuring OSI products exclusively in the very near future. The current understanding is that GOSIP should be used at RFP time in specifying future networking products. Other non-OSI products can also be acquired. To quote from the Draft FIPS document:

"This standard is effective __ (six months following publication). For a period of eighteen months after the effective date, agencies are permitted to acquire alternative protocols which provide equivalent functionality to the GOSIP protocols. Agencies are encouraged to use this standard for solicitation proposals for new network products and services to be acquired after the effective date. This standard is mandatory for use in all solicitation proposals for new network products and services to be acquired after __ (eighteen months after the effective date).

For the indefinite future, agencies will be permitted to buy network products in addition to those specified in GOSIP and its successor documents. Such products may include other non-proprietary protocols, proprietary protocols, and features and options of OSI protocols which are not included in GOSIP."

Internet Engineering Task Force reorganized

Coordinate efforts

In order to facilitate Internet technical coordination between various agencies (including DARPA, DCA, NSF, NASA and OSD), the Internet Engineering Task Force has joined forces with the NSFNET Routing Group.

Working groups

This new and very large combined technical group, is organized as a small central coordinating committee and a varying number of ad hoc and standing working groups. Future IETF meetings would be composed of detailed working group status reports, other technical presentations, and an opportunity for the working groups to meet.

Met jointly with ANSI

The Internet Engineering Task Force met in Boston on April 22-24. This was a very large meeting for two reasons: it was the first meeting since the IETF reorganization (described above), and it also included a joint session with the ANSI Network Layer Group X3S3.3. Since the ANSI group is working on several items of concern to the IETF (Routing, Congestion Avoidance, etc.), but with "ISO focus," it was felt that such a joint meeting was very helpful, particularly as the DoD considers migration to ISO. It is hoped that such joint meetings will take place in the future as well.

The High-Level Entity Management System (HEMS)

by Craig Partridge, NNSC at BBN Laboratories
and Glenn Trewitt, Stanford University

This article gives an overview of the High-Level Entity Management System. This system is experimental, and is currently being tested in portions of the Internet. It is hoped that this work will shortly lead to a standard for IP internetwork monitoring.

Used to be simple

Until recently, a majority of critical components in IP networks, such as gateways, have come from a very small set of vendors. While each vendor had their own set of management protocols and mechanisms, the collection was small, and a knowledgeable system administrator could be expected to learn them all.

Problem grows as Internet grows

Now, however, the number of vendors has grown quite large, and the lack of an accepted standard for management of network componenets is causing severe management problems. Compounding this problem is the explosive growth of the connected IP networks known as the Internet. The combination of increased size and heterogeneity is making internetwork management extremely difficult. This article discusses an effort to devise a standard protocol which should help alleviate the management problem.

A detailed description of the High-Level Entity Management System will be issued as a set of RFCs. This set is expected to change and grow over time, and readers are strongly encouraged to check the RFC Index to find the most current versions.

Monitoring and control

Historically the IP community has divided network management into two distinct types of activities: monitoring and control. *Monitoring* is the activity of extracting or collecting data from the network or a part of the network to observe its behavior. *Control* is the activity of taking actions to effect changes in the behavior of the network or a part of the network in real-time, typically in an attempt to improve the network's performance.

Note that the ability to control presupposes the ability to monitor. Changing the behavior of the network without being able to observe the effects of the changes is not useful. On the other hand, monitoring without control makes some sense. Simply understanding what is causing a network to misbehave can be useful.

Because effective monitoring is the key first step, the authors have concentrated their attention on defining an effective monitoring system. How the system might be used for control has received less attention, although possible mechanisms are discussed.

Overview of the HEMS

The HEMS is made up of three parts: a query processor which can reside on any entity which has an IP address, a trap generator which also resides on IP entities, and applications which know how to send requests to the query processor and interpret the replies. The query processor and applications communicate using a management protocol which runs over a standard transport protocol.

Query processor

The query processor is the key to the management system, since it is the source of all monitoring information and the processor of all control requests. For optimal network management, we would like to see query processors on many network entities.

To encourage the implementations of query processors, one of the primary goals in designing the query processor was to make it as small and simple as possible, consistent with management requirements and good performance.

Defining the management requirements was no small task, since the networking community has not yet reached a consensus about what kinds of monitoring information should be available from a network, nor what control functions are required to properly manage a network. The standards for HEMS were developed through communications with several interest groups, and represent the authors' best effort to distill the varying sets of needs.

Robust and extensible

The authors settled on a system which was extensible, robust and host-architecture-independent, and as simple as possible, consistent with the other goals. Extensibility was essential because it is clear that management needs will continue to evolve, and a closed system, which could not be changed, would be obsolete almost as soon as it was defined. Unfortunately, extensibility is also the requirement least consistent with simplicity, since the need to make the system extensible led the authors to use self-describing data formats and an interpreted query language. A robust system is required if the system is to be useful for diagnosing network failures. If the monitoring system cannot survive at least moderate network failures, it is not useful.

ISO ASN.1 encoding

The query processor is designed to be highly extensible. An application sends the query processor instructions about objects to be examined or changed. The query processor locates the values in its host entity, and performs the requested operations. The values are self-describing, using the binary-encoding scheme defined in ISO Standard ASN.1. Care has been taken to use a limited set of the ASN.1 coding set, so that query processor's handling of data can be optimized. It should not be necessary to support an ASN.1 data compiler in the query processor, though it may be required of the application.

The objects that a host entity maintains are defined on an entity type basis. Every participating IP entity is required to support a small set of objects that are required of all IP entities, plus a larger set of objects which are required of all entities of a given type (e.g., gateways, hosts, access machines). Entities are permitted to make additional, entity-specific objects available to applications. A method for discovering the existence of additional objects is defined. The combination of self-describing data, the ability to add to the standard data set and query language which can be easily enhanced appeared to offer the necessary extensibility.

Trap generator

On many network entities, particularly critical network components such as gateways, it is necessary to have a way for the devices to send unsolicited status messages to network management centers. In the IP community, these messages are called "traps."

In the HEMS system, traps are handled as slightly specialized replies to queries, and are sent to one or more management centers. Like all other HEMS messages traps are formatted in ASN.1 format.

HEMS (continued)

Applications	<p>The HEMS expects that applications will be more intelligent than the query processor. Among other functions, the applications will have to be able to identify and parse entity-specific values which may be returned.</p> <p>The details of applications are largely not discussed in the HEMS specifications because there is very little that needs to be standardized. Applications must send requests using the protocols discussed in the next section, but the interfaces the applications provide for displaying monitoring or control information is entirely application dependent.</p>
Protocols	<p>Query processors and applications communicate using an application-specific monitoring protocol, the High-Level Entity Management Protocol (HEMP). This protocol provides the formatting rules for the queries and their replies.</p> <p>HEMP runs over a standard transport protocol. There was a certain amount of debate in the community about what type of transport protocol was best suited for monitoring. The key issue was how reliable monitoring interactions needed to be.</p> <p>The authors expect that three types of management activities will predominate: status monitoring, firefighting and receipt of traps.</p>
Status monitoring	<p>Status monitoring is envisioned as occasional retrieval of monitoring information, possibly in response to the receipt of trap messages. In these situations, the network is expected to be in good working condition, and monitoring exchanges could probably comfortably work with an unreliable transport protocol. The chance of data loss is small, and probably not a serious problem since the data is unlikely to be so important that it must be reliably delivered. (However, it should be noted that some applications may prefer reliable delivery because it is more convenient).</p>
Firefighting	<p>Firefighting is a completely different situation. In this scenario, one or more sites are using management applications to try to locate and fix a network problem. Here we must assume that while the network functions (i.e. data can get through), it is not very healthy. We should assume that packets are being lost, that network routes will be non-optimal, and that it is essential that the monitoring data (which is presumably diagnostic) get back to the application and that control requests are reliably delivered to the entity. In such circumstances, a reliable protocol is essential.</p>
Traps	<p>Traps provide yet another bit of complexity. Traps contain useful status information, but experience suggests that this information does not have to be delivered reliably. If the problem is serious enough, it will re-occur and the trap will be sent again. Furthermore, traps will often be sent to more than one management center, which would appear to preclude the use of connection-oriented, reliable protocols such as TCP for traps.</p>
Transaction protocol	<p>The current decision has been to establish two possible transport options for HEMS. More experimental systems may use the Versatiles Message Transaction Protocol (VMTP), an experimental IP transaction protocol. Near term production systems can use a combination of the Transmission Control Protocol (TCP) and the User Datagram Protocol (UDP).</p>

University of Wisconsin implements ISO OSI protocols

An extensive set of networking software following the Open Systems Interconnection (OSI) standards has been implemented at the University of Wisconsin Madison. The project, supported by International Business Machines, has produced a complete networking environment under the UNIX operating system for the IBM RT/PC workstation. Included in the package are both Connection-less and Connection-Oriented Network Services, Transport classes 0 and 4, Session, and Presentation protocols. In addition, the package includes Message Handling Service (MHS) and File Transfer, Access, and Management (FTAM) applications. Professors David DeWitt, Lawrence Landweber, and Marvin Solomon are supervising the project. The implementation team includes Rob Hagens, Nancy Hall, Sue Lebeck, and Derek Zahn.

Implementation challenge

The major challenge of this project was working from networking standards that are relatively young, and therefore largely unused and untested. The immaturity of the protocol and service definitions and the dearth of implementation experience contribute to the challenge of implementing OSI. Indeed, even existing implementations face interoperability difficulties, though much progress has been made. Compatibility has been enhanced by numerous implementors' agreements developed by such groups as the participants in the NBS OSI Implementors' Workshops, the Standards Promotion and Application Group (SPAG), and the Committee for European Normalization (Electronique) (CEN-CENELEC).

Extensibility can be a problem

In some areas, particularly in the higher layers, the OSI architecture and protocols demonstrate a clear concern for extensibility. Examples include Presentation's provision for user-defined presentation contexts, and MHS' provision for additional message body types. Many of the service specifications define functional subsets of the layers' services to accommodate various categories of potential OSI users. Unfortunately, this variety of functional subsets can adversely affect interoperability.

In other areas, a certain short-sightedness and unaccommodating attitude tend to prevail. Examples include backwards class negotiation in the transport layer, and a lack of elegance in the network layer for accommodating diverse subnetworks.

Reason for optimism

Despite these difficulties, there is reason for optimism. Any international endeavor encompassing such a wide scope and demanding such extensive cooperation can be expected to find its path to success a rocky one. The architects and implementors of OSI have already removed many boulders and traveled a long way. As OSI products emerge and come into widespread use, the hope is that the path of interconnection will be transformed into a smooth pavement.

Profile: NSFNET

When the National Science Foundation (NSF) established its national supercomputer centers in 1985, it also planned to create a communications network that would give remote locations access to these state-of-the-art facilities. NSF planners envisioned a system they dubbed "NSFNET." Based on a "backbone" connecting the supercomputer centers, NSFNET would combine existing networks and newly created ones into an internet, or network of networks, to serve the centers and their users. In addition to gaining access to the centers' computing technology, researchers at geographically dispersed locations would be part of a nationwide research network across which they could exchange scientific information.

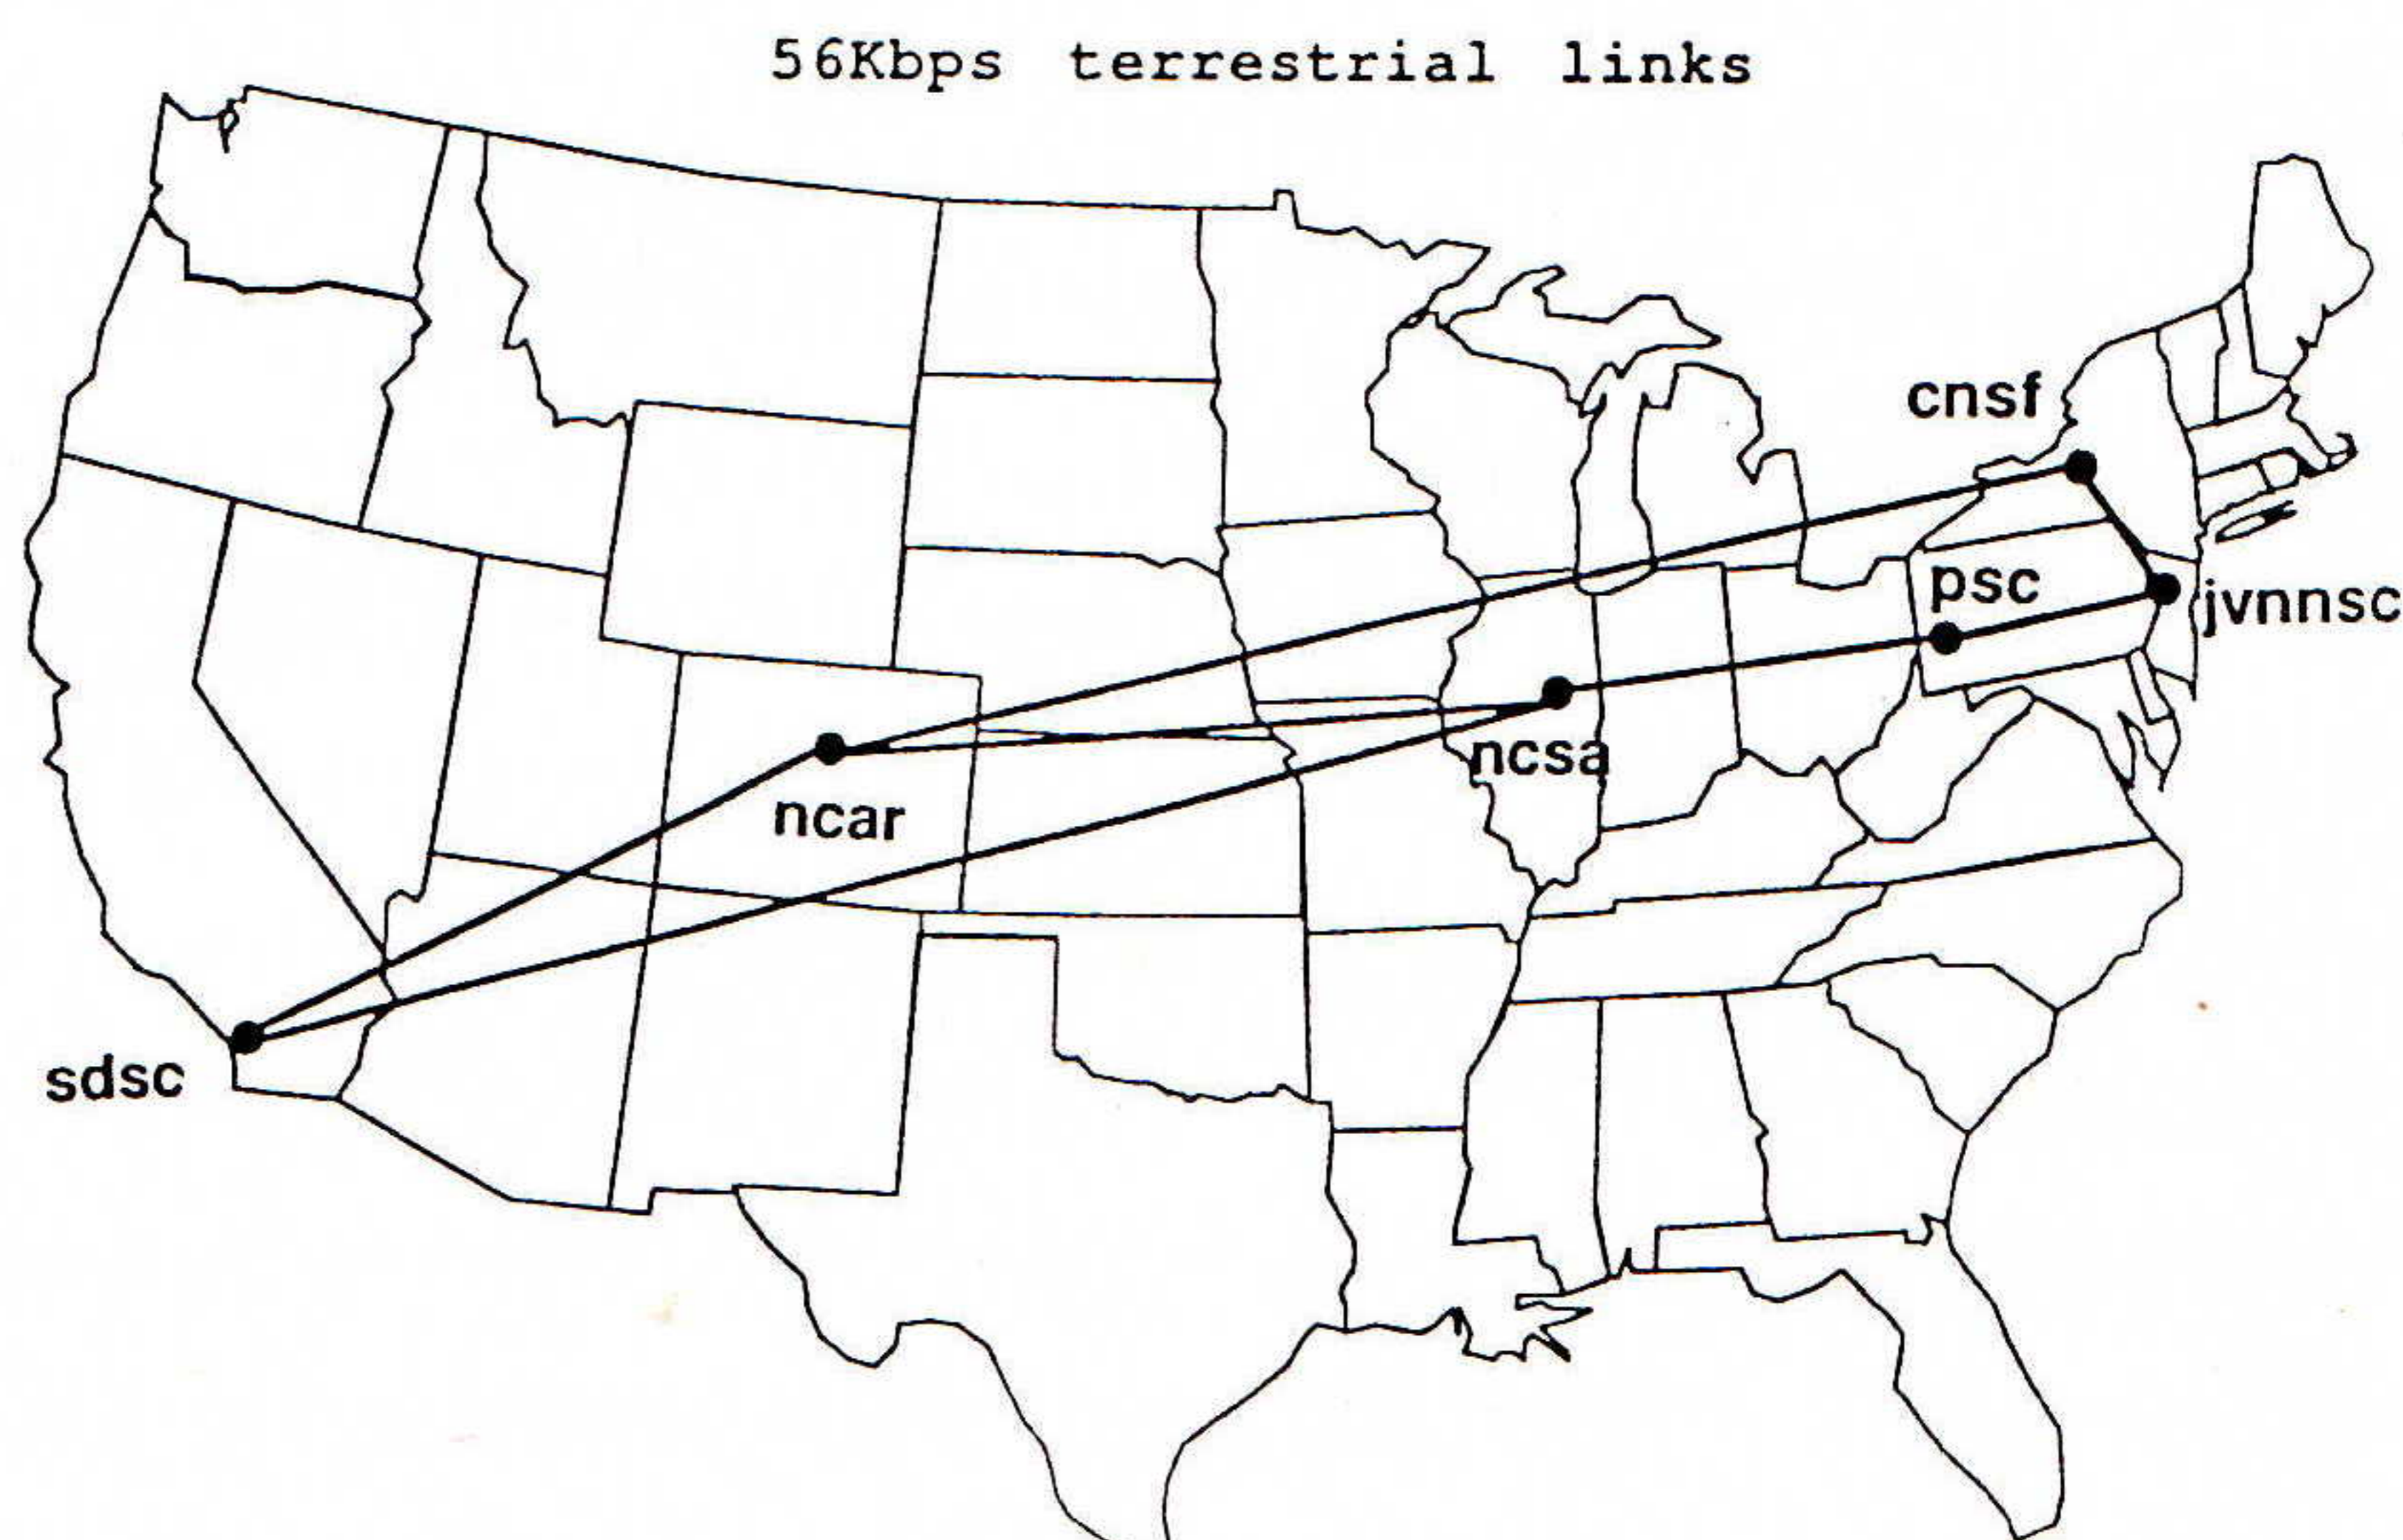
Access to resources and information sharing

"Although the primary role of NSFNET remains access to NSF-funded supercomputers and other unique scientific resources, its use as a general-purpose network, which enables scientists to share research findings, is becoming increasingly important," says Stephen Wolff, Director of NSF's new Division of Network and Communications Research and Infrastructure, part of the Computer and Information Science and Engineering Directorate.

NSFNET components

NSFNET is organized as a three-level hierarchy: the backbone; autonomously-administered wide-area networks serving communities of researchers; and campus networks. The backbone has been in use since July 1986 and is fully operational. It provides redundant paths among NSF supercomputer centers. While several wide-area networks are already connected to the NSFNET backbone, more are being built with partial funding from NSF and will be connected as they are completed. Based on the number of actual connections and those in the planning stages, Mr. Wolff expects that 150 or more university sites will be on the network by the end of 1987.

NSF Backbone Network



**Supercomputer
centers**

NSF created the supercomputer centers in response to a growing concern that a lack of access to sophisticated computing facilities had severely constrained academic research. A project solicitation in June 1984 resulted in the creation of the following centers:

- John von Neumann National Supercomputer Center, Princeton, New Jersey (JVNNSC)
- San Diego Supercomputer Center, University of California (SDSC)
- National Center for Supercomputing Applications, University of Illinois (NCSA)
- Cornell National Supercomputer Facility, Cornell University (CNSF)
- Pittsburgh Supercomputer Center, operated jointly by Westinghouse Electric Corp., Carnegie-Mellon University, and the University of Pittsburgh (PSC)
- The Scientific Computing Division of the National Center for Atmospheric Research, Boulder, Colorado (NCAR)

All the centers are multidisciplinary and are available to any researcher who is eligible for NSF support. They offer access to computers made by Cray Research Inc., Control Data Corporation, ETA and IBM.

Protocols

NSFNET is using the TCP/IP protocols of the DARPA Internet as the initial standard. The system will work toward adopting international standards as they become established. The protocols link networks that are based on different technologies and connection protocols, and provide a unified set of transport and application protocols. As the NSFNET system continues to evolve, the typical user working at a terminal or workstation will be able to connect to and use various computer resources -- including the supercomputer centers -- to run interactive and batch jobs, receive output, transfer files, and communicate with colleagues throughout the nation via electronic mail. Most researchers will have either a terminal linked to a local super-minicomputer or a graphics workstation. These will be connected to a local area network that is connected to a campus network, and, via a gateway system, to a wide-area network.

Management

Four institutions are sharing the interim management of NSFNET: the University of Illinois (overall project management and network engineering), Cornell University (network operations and initial technical support), University of Southern California Information Sciences Institute (protocol enhancement and high-level technical support), and the University Corporation for Atmospheric Research (management of the NSF Network Service Center through a contract with Bolt Beranek and Newman Laboratories Inc.). The long-term management and operational structures of NSFNET are under study, and a permanent arrangement should be in place in 1987.

continued on next page

Profile: NSFNET (*continued*)

NSF Network Service Center

The NSF Network Service Center (NNSC) is providing general information about NSFNET, including the status of NSF-supported component networks and supercomputer centers. The NNSC, located at BBN Laboratories Inc., in Cambridge, MA, is an NSF-sponsored project of the University Corporation for Atmospheric Research.

The NNSC, which currently has information and documents on-line and in printed form, plans to distribute news through network mailing lists, bulletins, newsletters, and on-line reports. The NNSC also maintains a database of contact points and sources of additional information about the NSFNET component networks and supercomputer centers.

When prospective or current users do not know whom to call concerning their questions about NSFNET use, they should contact the NNSC. The NNSC will answer general questions, and, for detailed information relating to specific components of NSFNET, will help users find the appropriate contact for further assistance. In addition the NNSC will encourage the development and identification of local campus network technical support to better serve NSFNET users in the future.

Users may reach the NNSC by calling the hotline 617-497-3400 or by sending electronic mail to nnsc@nnsc.nsf.net. Karen Roubicek, (roubicek@nnsc.nsf.net) is the NNSC User Liaison.

Connecting to NSFNET

NSFNET is part of a collection of interconnected IP networks referred to as the Internet. IP, the Internet Protocol, is a network protocol which allows heterogeneous networks to combine into a single virtual network. TCP, the Transmission Control Protocol, is a transport protocol which implements the packet loss and error detection mechanisms required to maintain a reliable connection between heterogeneous computers on diverse networks. An example of an application which uses TCP/IP is TELNET, which provides virtual terminal service across the network.

Only IP-based networks can connect to the Internet; therefore an organization that plans to use NSFNET either must have an existing IP network or have access to one. Many large universities and technical firms already have links to the Internet in place. The computer science department of a university or the engineering support division of a company are most likely to have IP connectivity or to have information on the local connections that exist. Prospective users can ask the NNSC to determine whether an organization is already on the Internet.

If an organization does not have an IP link, it can obtain one in several ways:

- NSF has a program that funds the connecting of organizations to the NSF regional/state/community networks that are part of NSFNET. The NNSC has more information on this program.

- The Computer Science Network, CSNET, provides gateway service to several IP networks, including NSFNET. To get CSNET service an organization must become a CSNET member. For more information call or write the CSNET Coordination and Information Center (CIC), BBN Laboratories Inc., 10 Moulton Street, Cambridge, MA 02238, 617-497-2777. The CIC's network address is `cic@sh.cs.net`.
- Users may be able to get access to NSFNET through time-share accounts at other organizations such as local universities or companies.

Some supercomputer centers support access systems other than NSFNET, such as BITNET, commercial X.25 networks, and dial-up lines, which do not use IP-based protocols. The Supercomputer Centers' user service organizations can provide more information on these alternatives.

The following is a list of the NSFNET component networks:

**State and Regional
Networks**

BARRNET California's Bay Area Regional Research Network
MERIT Michigan Educational Research Network
MIDNET Midwest Network
NYSERNET New York State Educational and Research Network
SESQUINET Texas Sesquicentennial Network
SURANET Southeastern Universities Research Association Network
WESTNET Southwestern states Network
NORTHWESTNET Northwestern states Network

**Consortium
Networks**

JVNCNET connects the John von Neumann National Supercomputer Center at Princeton, NJ, with a number of universities.

PSCAANET is the network of the Pittsburgh Supercomputer Center Academic Affiliates group.

SDSCNET is centered at the San Diego Supercomputer Center.

**Other wide-area
networks**

ARPANET developed by the Defense Advanced Research Projects Agency of the US Department of Defense, the Arpanet has over 150 sites and will expand as new sites are added for use by the NSFNET community.

BITNET links computers at 175 universities across the nation, generally connecting campus computing centers.

CSNET has members at 200 sites from universities, businesses, government and nonprofit agencies. University members typically are computer science departments.

UNIDATA is UCAR's planned system to link the atmospheric sciences community. UNIDATA, which will be an important component of the NSFNET configuration for atmospheric researchers, should be operational by 1988 and will tie into NSFNET, as well as provide weather data and other services.

Of Waves and Whitecaps

by M. A. Padlipsky

Only a few close friends (and perhaps a very few very close readers of The Book) are aware of it, but I started out as one of TCP's critics. Since I can't assume that everybody who reads this even knows who I am-- much less what The Book is, but I'll leave it up to the Editor as to whether it gets plugged here at all, and where it does if it does-- perhaps it ought to be observed that that's actually a euphemism for "one of TCP's most outspoken, maybe even strident, critics." Yet somehow, from that starting point I seem to have evolved into a sort of Auxilliary Bishop of the Congregation for the Propagation of the Faith, and it might be of particular interest to the readers of this particular journal to learn how that came about. (Anyway, Ole asked me for a few words on an appropriate topic of my own choosing, and that's what I decided they'd [nominally] be about; if you don't care, I don't care.)

Back when TCP was being touted as the Wave of the Future-- say, ten years or so ago-- it actually did have quite a few flaws. To oversimplify: there weren't many implementations around-- maybe half a dozen-- and they neither operated nor interoperated very well; the spec wasn't very clear and it contained some very weird notions (if you've never heard of Rubber EOLs, be grateful); but the major problem, as far as I was concerned, was that you had to have all those bloody "reliability" mechanisms in play on every transmission, whether you liked it or not (i.e., IP hadn't been separated from TCP yet). Of course, I was also decidedly unamused by the fact that Vint was being so stringent in his application of his "only implementors get to design it" rule that I was deriving all these impressions only at second hand, but that's special pleading. At any rate, I made myself a little sign that said "TCP: The Whitecap of the Future," put it up on a wall of my office, and forgot about it.

So what happened to turn the Saul-analogue into a Paul-analogue? Two things happened, actually. In the first place, TCP grew up. Implementations increased in number and speed, and incompatibilities decreased in number and graveness; the spec [*see pg 14] was iterated, and improved (which doesn't automagically follow); and mainly, to my mind at least, they finally "got the Layering right" and TCP and IP were defined as separate entities.

Also-- and this is a biggie, but one so easy to overlook that it didn't even get mentioned until the second draft-- TCP/IP "fit in" quite naturally with an entire "suite" of protocols; that is, the Host-Host protocol is far from the whole story: Telnet, and FTP, and Mail and the like matter even more than TCP versus "NCP" when you get down to using intercomputer networks instead of philosophizing about protocols, and the important thing is that the ARPANET already had them, and fairly rapidly got them integrated with TCP/IP since the latter had, after all, been designed by people who were intimately familiar with the "Process/Applications Layer" protocols and hence didn't introduce disparities very often (if at all). In the second place, to overstate it, I threw up and then I grew up. What I mean is that in the meantime I became aware of another touted Wave of the Future (we're up to around 1980 by now).

I looked upon it, and found that it was so bad that I might even have opted for TCP over it even if they hadn't gotten the layering right, since TCP was at least being designed by implementors even if I wasn't one of them, whereas the other would-be Wave was being concocted by standards committee members. Now, this doesn't strike me as the right context to be particularly Constructively Snotty about that other soi-disant Wave, a/k/a the ISORM (which some call the International Standards Organization's Open System Interconnection Reference Model, others the Pestilential Palindrome, still others still worse things, but, sadly, which all too many would be calling "my ricebowl" if they were intellectually honest). After all, if you're a TCP/IPer reading this you probably already know your enemy, and if you're an ISORMite reading this in order to know your enemy I might applaud the uncharacteristically scholarly gesture you're making, but why should I save you the purchase price of The Book to find out about how your ricebowl is misgilded? (One recent development does bear note, however: I find it absolutely hilarious that the ISORM Security Appendix Committee is planning to vest no Security functionality in the "Session Layer". The Why's & Wherefores would take us too far afield, but do think about it.) Suffice it to say, though, that I meant it when I said that TCP/IP could have struck me as being far worse on an absolute scale and I'd still have found the ARPANET Suite better on a relative scale.

Sooooooo... having concluded that the ISORM Suite was an Eternal Whitecap, and having realized that there really were only two choices if the object of the game were heterogenous intercomputer networking (sorry, XNS, and sucks to you, SNA), it wasn't all that hard to get to the point where I can (pseudo)conclude this little guest editorial or sermon or pep-talk or whatever it's turning out to be with the following:

Presumed typical reader of *ConneXions*, I salute you: you're riding the Wave of the Present, and if you do it right that wave may never become a whitecap, since the wave allegedly coming along behind it will probably always be too far out to sea to ride, and even if it ever does get somewhere it'll probably turn out to be in another ocean entirely, anyway.

Well, OK, I guess that's too metaphorical/cryptic even for me. What I meant it to mean is something like this: There's still meaningful work to be done in the TCP/IP "world" (including the Process/Applications Layer protocols), ranging from resolving misinterpretations ("4.2"'s reportedly highly idiosyncratic view of how to deal with multiple Internet Addresses comes to mind as just one example) to furnishing needed new functions (dealing appropriately with "congestion," e.g.).

Not actively pursuing such matters on grounds of "only getting it to work and not worrying about getting it to work well because the ISO will be along any year now and we're just doing a job here" would, I submit, be an entirely wrong way to view the situation.

continued on next page

Of Waves and Whitecaps *(continued)*

Look at it this way: In ten years, we've gone from a few limping implementations to/of a somewhat muddled conceptual model to a few dozen (or a few score, who keeps count?) respectable-to-gutsy implementations to/of a rational conceptual model, while the Other Side still hasn't even admitted its "reference model" is fundamentally muddled. Sure sounds like we're still way ahead to me, and we might as well stay there.

Look at it another way: it took around nine years for us to mature to the point where we realized that commercial implementors needed a forum in which to clean up the relatively easy loose ends with help from the designers, while the Other Side doesn't really seem even to have designers to have recourse to, they've just got standards committee members (who, apparently not understanding their own definition of Layering in great enough numbers, keep putting L6 sorts of things in L7, to pick on just one of my favorite ISOlecisms, again without bothering with the W&W's). Sure sounds like we'll probably always be way ahead to me.

So even if it weren't for the Padlipsky's Law which holds that implemented protocols have barely finite intertias of rest, I still wouldn't think that it would be appropriate to view the continuing work on the ARPANET Protocol Suite as "just going through the motions until the real thing comes along" (though it should be clear that I have some fear that some readers of this might be inclined to view it in just that fashion or I wouldn't be bothering to attempt to gloss the "concluding" metaphor), because I expect the ISO stuff to be sufficiently bad when/if it arrives that even the DoD will have to come to realize that it's foolish to junk a well-maintained late-model Buick just because they're committed to throw bad money after good and buy a Renault for use on some particular trips.

Hmmmmmm. I seem to have glossed one metaphor with several more metaphors. Ah, well, this is only a Guest Whatever, after all. And it's probably gone on too long as it is, so I'd better hurry up and get to the non- (because it's meant to be part of the paper, actually) footnote and then get off stage/page:

[*] Note extremely well that the "improved spec" I was talking about is Postel's RFC 791/3, not the bloody "Mil-Stds," which, if they had been what they were supposed to be, would have obviated the need for an implementors' forum so you wouldn't even have had the opportunity to misunderstand or ignore all this in the first place. (Indeed, it's the apparent impossibility of doing a formal spec right that makes me fairly confident that the ISOids will always be trying to ride a whitecap.)

Not to be coy about it, but I'm just not up for picking on the Mil-Stds in any detail here; however, as a friendly hint to the presumed typical reader of this thing, if you've only looked at the Mil-Stds thus far, do go look at the RFCs (or their latest successors): they do a much better job on the W&W's, and to my taste are even better on the Whats too, since I'm allergic to state machinery. (As another friendly hint, think very hard about the fact that the interface to an interpreter of a protocol is NOT the protocol.) And if you've got any pull with DCA, try to convince them to get my old redaction of the RFCs resurrected, typed up, and made available: I think it does a somewhat better job still on the W&W's.

(And in case you haven't been paying attention, I'll throw in a final friendly gesture and remind you that W&W's means Whys and Wherefores and that without understanding them I don't think you can do the implementations right-- a desire to do which is, I hope, the underlying reason for your reading this journal in the first place-- but if you don't know what a redaction is and are too lazy to find out, then clearly you don't deserve my help, so I'll stop right here, right now.)

May you all enjoy the ride on the wave, and get to stay away from the whitecap.

Cheers,

MAP

(Institutional affiliation withheld as usual, to avoid necessity of Corporate Review.)

Michael A. Padlipsky is one of the "Network Old Boys." He participated in the design of the TCP/IP protocol suite and has written a number of RFCs. He is also known for The Book: "The Elements of Networking Style," published by Prentice-Hall.

UNIX is a trademark of AT&T Bell Laboratories.

MS is a trademark of Microsoft Corporation.

VMS is a trademark of Digital Equipment Corporation.

IBM is a trademark of International Business Machines Corporation.

Xenix is a trademark of Microsoft Corporation.

CONNEXIONS

480 San Antonio Road
Suite 100
Mountain View, CA 94040

FIRST CLASS MAIL
U.S. POSTAGE
PAID
SAN JOSE, CA
PERMIT NO. 1

CONNEXIONS

PUBLISHER Daniel C. Lynch

EDITOR Ole J. Jacobsen

EDITORIAL ADVISORY BOARD Dr. Vinton G. Cerf, Vice President, National Research Initiatives.

Dr. David D. Clark, The Internet Architect, Massachusetts Institute of Technology.

Dr. David L. Mills, NSFnet Technical Advisor; Professor, University of Delaware.

Dr. Jonathan B. Postel, Assistant Internet Architect, Internet Activities Board; Division Director, University of Southern California Information Sciences Institute.

Subscribe to CONNEXIONS

U.S./Canada \$100. for 12 issues/year \$180. for 24 issues/two years \$240. for 36 issues/three years
International \$ 50. additional per year (Please apply to all of the above.)

Name _____ Title _____

Company _____

Address _____

City _____ State _____ Zip _____

Country _____ Telephone () _____

☐ Check enclosed (in U.S. dollars made payable to CONNEXIONS).

☐ Charge my ☐ Visa ☐ Master Card Card # _____ Exp. Date _____

Signature _____

Please return this application with payment to:

CONNEXIONS

480 San Antonio Road Suite 100
Mountain View, CA 94040
415-941-3399

Back issues available upon request \$10./each
Volume discounts available upon request

CONNEXIONS